

Variational inference for large-scale models of discrete choice

Michael Braun

MIT Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02142

Jon McAuliffe

Statistics Department
University of Pennsylvania
Wharton School
Philadelphia, PA 19104

6 January 2008*

Abstract

Discrete choice models are commonly used by applied statisticians in numerous fields, such as marketing, economics, finance, and operations research. When agents in discrete choice models are assumed to have differing preferences, exact inference is often intractable. Markov chain Monte Carlo techniques make approximate inference possible, but the computational cost is prohibitive on the large data sets now becoming routinely available. Variational methods provide a deterministic alternative for approximation of the posterior distribution. We derive variational procedures for empirical Bayes and fully Bayesian inference in the mixed multinomial logit model of discrete choice. The algorithms require only that we solve a sequence of unconstrained optimization problems, which are shown to be convex. Extensive simulations demonstrate that variational methods achieve accuracy competitive with Markov chain Monte Carlo, at a small fraction of the computational cost. Thus, variational methods permit inferences on data sets that otherwise could not be analyzed without bias-inducing modifications to the underlying model.

1 Introduction

Discrete choice models have a long history in statistical analysis, appearing in applications as varied as the analysis of consumer choice data (Guadagni and Little 1983; Fader and Hardie 1996), transportation planning (Theil 1969; McFadden 1974; Ben-Akiva and Lerman 1985), economic demand estimation (Train et al. 1987; Revelt and Train 1998), new product development (Moore et al. 1999), portfolio analysis (Uhler and Cragg 1971) and

*The authors gratefully acknowledge research assistance from Sandy Spicer and Liz Theurer.

health services deployment (Hall et al. 2002). They apply to situations where agents (also called choosers or decision-makers) select items from a finite collection of alternatives (the choice set), either once or repeatedly over time. For example, in a marketing context, agents are “households”; each household makes a number of “trips” to a store, and we observe the items selected for purchase on each trip.

Heterogeneous discrete choice models, which allow preferences to differ across agents, are based on a hierarchical regression formulation. We have agents numbered $h = 1, \dots, H$, each with an unseen parameter vector θ_h encoding preferences over item attributes. We observe one or more choice events $Y_h \sim p(y_h | \theta_h)$ per agent. The θ_h ’s are modeled as independent draws from a prior distribution $p(\theta_h | \phi)$, where ϕ is a hyperparameter. This prior represents the heterogeneity of preferences across the population. Inference in such a hierarchical model allows us to pool information across decision-makers. If we use an empirical Bayes point estimate of ϕ (Robbins 1955), the posterior distribution of each θ_h depends on all of Y_1, \dots, Y_H , through the common estimate $\hat{\phi}$. In a fully Bayesian setup, integrating out the random variable ϕ creates similar dependence.

The marginal likelihood corresponding to one agent in a heterogeneous model is

$$p(y_h | \phi) = \int p(y_h | \theta_h) p(\theta_h | \phi) d\theta_h . \quad (1)$$

In most cases, including the “random utility” discrete choice model we study in this paper, (1) does not exist in closed form. As a consequence, we must use approximate methods both for empirical Bayes and fully Bayesian inference.

A standard empirical Bayes technique is to approximate (1) using Monte Carlo integration. But to match the asymptotics of maximum likelihood, the number of draws per agent must grow faster than the square root of the number of agents (Train 2003), which is infeasible for large-scale problems. The usual approach to the fully Bayesian random utility model is Markov chain Monte Carlo (MCMC) simulation (Albert and Chib 1993;

Allenby and Lenk 1994; Allenby and Rossi 1999). MCMC provides approximate draws from the joint posterior distribution on $\theta_1, \dots, \theta_H$ and ϕ . The draws enable the estimation of (1) and related integrals. However, the more agents there are in the data set, the more MCMC output we need to collect and store – even if we are only interested in ϕ , we still need repeated draws for all of $\theta_1, \dots, \theta_H$.

Variational methods (Jordan et al. 1999; Wainwright and Jordan 2003) offer a deterministic alternative for approximate inference. With variational inference, we maximize a data-dependent lower bound on the marginal likelihood (1), over a set of auxiliary parameters distinct from the model parameters. In the fully Bayesian specification, the end result is an approximate joint posterior distribution for $\theta_1, \dots, \theta_H$ and ϕ . For empirical Bayes, variational techniques lead to a point estimate $\hat{\phi}$ as well as an approximate posterior distribution for the θ_h 's.

The main advantage of variational methods versus MCMC is computational efficiency. Variational inference algorithms typically converge to their final approximation in far less time than it takes to generate an adequate number of MCMC draws. This advantage comes at the cost of a biased approximation, in contrast to the consistency guarantees that accompany MCMC. We give evidence in Section 4 that, for our random utility discrete choice model, variational bias is negligible, and the computational speedup is very large. Variational convergence is also easy to assess, in contrast to MCMC.

Furthermore, the size of the variational representation is fixed, while the size of the MCMC approximation increases with the number of draws. Variational techniques can therefore be applied to much larger data sets. For example, with 10,000 decision-makers and 20,000 total MCMC draws, using a 25-dimensional θ_h and corresponding ϕ , the MCMC representation of the posterior exceeds 2 GB—if we discard half the draws for burn-in of the chain. In fact, many data sets today contain observations on millions of agents, models can contain far more than 25 preference parameters, and MCMC chains may require

hundreds of thousands of iterations.

These difficulties are well known. MCMC is rarely applied to large-scale heterogeneous models. Indeed, to address scalability, it is common to work with data from a subset of individuals, or a subset of choice items. However, this approach discards information that is valuable in the inferential process, and it can lead to biased estimates (Bradlow and Zanutto 2006).

In this paper, we derive variational algorithms for a common discrete choice model – the mixed multinomial logit (MML) model. We study this model because it the workhorse of discrete choice theory and is well known in many disciplines, including economics and marketing. There are other popular discrete choice models in the literature, but the mixed multinomial logit has appeal: it is conceptually simple, yet still exhibits the MCMC inference issues just described.

The rest of the paper is organized as follows. Section 2 presents the details of the MML model. In Section 3, we describe variational procedures suitable for empirical Bayes and fully Bayesian inference under the MML model. These procedures include a novel application of the delta method for moments to variational inference. In Section 4, we compare variational methods to MCMC for the MML model, using an extensive suite of simulated data sets. Section 5 closes with discussion and future directions. Technical arguments and derivations are relegated to Appendixes A, B, and C.

2 The mixed multinomial logit model of discrete choice

Let there be H agents, indexed $h = 1, \dots, H$. We observe a total of T_h choice-event outcomes for agent h . At each choice event, the agent selects from among a fixed set of J items, indexed $j = 1, \dots, J$. The items are differentiated according to K attributes, indexed $k = 1, \dots, K$. The j th item’s value for the k th attribute can vary across agents and from one choice event to another. For example, households might shop at different stores charging various prices for the same good, and the price of a good may change over

time within a single store. We denote by x_{ht} the $J \times K$ matrix of attribute values, also called covariates, that agent h encounters at her t th choice event. The j th row of x_{ht} is denoted x_{htj}^\top . The outcome of this choice event is the observed categorical random variable y_{ht} , which we represent as a $J \times 1$ indicator vector.

We use the observed (x_{ht}, y_{ht}) pairs to infer which attributes have the strongest association with item choice. To this end, let U_{htj} denote the utility that accrues to agent h if she chooses item j at her t th choice event. This approach, called a “random utility model” (Train 2003), assumes utility is a noisy linear function of the attributes: $U_{htj} = \beta_h^\top x_{htj} + e_{htj}$. Here, β_h is a $K \times 1$ vector of agent-specific “tastes” or “preference loadings” for the item attributes, and e_{htj} is a random error term representing unobserved utility.

We assume that each agent, at each choice event, selects the item maximizing her utility. In the mixed multinomial logit model, we further assume the random error terms e_{htj} are iid from a Gumbel Type 2 distribution. The implied choice probabilities turn out to be

$$P(y_{htj} = 1 \mid x_{ht}, \beta_h) = \frac{\exp(\beta_h^\top x_{htj})}{\sum_{j'} \exp(\beta_h^\top x_{htj'})}, \quad j = 1, \dots, J \quad (2)$$

(McFadden 1974). In discrete choice modeling, the right-hand side of (2) is called the “multinomial logit” distribution, denoted $\text{MNL}(x_{ht}, \beta_h)$. It is essentially the same as the multi-logistic function used in polychotomous logistic regression, and it is often called the soft-max function in machine learning research.

We further assume that $\beta_{1:H}$ are iid from a K -variate normal distribution with mean vector ζ and covariance matrix Ω , which we write as $\mathcal{N}_K(\zeta, \Omega)$. For empirical Bayes estimation, the model is now completely specified:

$$y_{ht} \mid x_{ht}, \beta_h \stackrel{\text{ind}}{\sim} \text{MNL}(x_{ht}, \beta_h), \quad h = 1, \dots, H, \quad t = 1, \dots, T_h, \quad (3)$$

$$\beta_h \mid \zeta, \Omega \stackrel{\text{iid}}{\sim} \mathcal{N}_K(\zeta, \Omega), \quad h = 1, \dots, H. \quad (4)$$

The top-level parameters ζ and Ω , to be estimated by maximum marginal likelihood, represent the distribution of attribute preferences across the population. In particular, Ω gives us information about the correlation of preferences between agents.

A fully Bayesian approach requires hyperprior distributions for ζ and Ω . As is standard, we use conditionally conjugate distributions:

$$\zeta \mid \beta_0, \Omega_0 \sim \mathcal{N}_K(\beta_0, \Omega_0), \quad \Omega \mid S, \nu \sim \mathcal{W}^{-1}(S^{-1}, \nu). \quad (5)$$

In (5), β_0 and Ω_0 are pre-specified hyperparameters; $\mathcal{W}^{-1}(S^{-1}, \nu)$ is the inverse Wishart distribution with scale matrix S^{-1} and ν degrees of freedom; and S and ν are hyperparameters fixed in advance. We call the fully Bayesian approach to MML model inference “hierarchical Bayes.”

3 Variational inference for the MML model

We have presented the component hierarchical distributions in the mixed multinomial logit model. Now we turn to the question of estimation and inference procedures. In the following, variable names inside of $p(\cdot)$ are used to distinguish among densities: we denote the pdfs in (3)–(5) by $p(y_{ht} \mid x_{ht}, \beta_h)$, $p(\beta_h \mid \zeta, \Omega)$, $p(\zeta \mid \beta_0, \Omega_0)$, and $p(\Omega \mid S, \nu)$, respectively. We let $\mathcal{D} = \{x_{ht}, y_{ht}\}$ denote all observed variables, i.e., the data.

For the empirical Bayes version of the MML model, the posterior density of the latent preference vectors, $p(\beta_{1:H} \mid \mathcal{D}, \zeta, \Omega)$, is

$$\prod_{h=1}^H \frac{p(\beta_h \mid \zeta, \Omega) \prod_{t=1}^{T_h} p(y_{ht} \mid x_{ht}, \beta_h)}{\int p(\beta_h \mid \zeta, \Omega) \prod_{t=1}^{T_h} p(y_{ht} \mid x_{ht}, \beta_h) d\beta_h}. \quad (6)$$

The joint posterior density for hierarchical Bayes, $p(\beta_{1:H}, \zeta, \Omega \mid \mathcal{D})$, is

$$\frac{p(\zeta) p(\Omega) \prod_{h=1}^H p(\beta_h \mid \zeta, \Omega) \prod_{t=1}^{T_h} p(y_{ht} \mid x_{ht}, \beta_h)}{\int p(\zeta) p(\Omega) \prod_{h=1}^H \int p(\beta_h \mid \zeta, \Omega) \prod_{t=1}^{T_h} p(y_{ht} \mid x_{ht}, \beta_h) d\beta_h d\zeta d\Omega}. \quad (7)$$

The numerator in both cases is the joint density of latent and observed variables, computed by multiplying together the densities defined in the model hierarchy (3)–(5).

The integrals appearing in these posterior densities have no closed form. As a consequence, exact inference is intractable. Variational inference is a deterministic alternative to the MCMC methods usually applied to this problem (Rossi et al. 2005). A variational algorithm selects from a pre-specified family of distributions \mathcal{Q} the best approximation to the true posterior distribution. We define \mathcal{Q} so that all of its members permit tractable probability calculations. Then, wherever we need the true posterior, such as for an expectation, we use the approximating variational distribution instead. This plug-in idea underlies MCMC methods as well – in place of the true posterior, we substitute the empirical distribution of the MCMC posterior draws.

3.1 Variational empirical Bayes

In this section we give an overview of a variational algorithm for approximate empirical Bayes estimation in the MML model. Appendix A fills in the details. We first specify a family of approximating distributions $\mathcal{Q} := \{q(\beta_{1:H} | \lambda) : \lambda \in \Lambda\}$ for the true posterior distribution (6). Since this posterior factors over h , we take \mathcal{Q} to be a family of factored distributions as well, so that $q(\beta_{1:H} | \lambda) := \prod_h q(\beta_h | \lambda_h)$. In particular, each factor $q(\beta_h | \lambda_h)$ is a K -variate normal density, with mean μ_h and covariance matrix Σ_h .

For the particular data set at hand, we want to find $q(\beta_{1:H} | \lambda^*)$, the best approximation in \mathcal{Q} to the posterior distribution $p(\beta_{1:H} | \mathcal{D}, \zeta, \Omega)$. To make the idea of a best approximation precise, we measure discrepancy with the Kullback-Leibler (KL) divergence (also called the relative entropy). Shortening $q(\beta_{1:H} | \lambda)$ to q_λ , the optimal variational parameters are given by

$$\lambda^* = \operatorname{argmin}_{\lambda \in \Lambda} \operatorname{KL}[q_\lambda || p]. \quad (8)$$

We can express the KL divergence between q_λ and p as

$$\text{KL}[q_\lambda \parallel p] = \mathbb{E}_{q_\lambda} \log \left[\frac{q(\beta_{1:H} \mid \lambda)}{p(\beta_{1:H} \mid \mathcal{D}, \zeta, \Omega)} \right] \quad (9)$$

$$= -\mathcal{L}(\lambda; \zeta, \Omega) + \log p(\mathcal{D} \mid \zeta, \Omega) \quad (10)$$

where we define

$$\mathcal{L}(\lambda; \zeta, \Omega) := -\mathbb{E}_{q_\lambda} \log \left[\frac{q(\beta_{1:H} \mid \lambda)}{p(\beta_{1:H}, \mathcal{D} \mid \zeta, \Omega)} \right]. \quad (11)$$

Here, \mathbb{E}_{q_λ} denotes an average over $\beta_{1:H}$, using $q(\beta_{1:H} \mid \lambda)$. Because $\text{KL}[q_\lambda \parallel p]$ is non-negative, (10) implies that, for all distributions q_λ ,

$$\mathcal{L}(\lambda; \zeta, \Omega) \leq \log p(\mathcal{D} \mid \zeta, \Omega). \quad (12)$$

The likelihood is called the “evidence” in some contexts; we therefore christen $\mathcal{L}(\lambda; \zeta, \Omega)$ the *evidence lower bound (ELBO)* function.

Using (10), we can formulate a maximization problem equivalent to (8), having the same optimal point $q(\beta_{1:H} \mid \lambda^*)$:

$$\lambda^* = \underset{\lambda \in \Lambda}{\operatorname{argmax}} \mathcal{L}(\lambda; \zeta, \Omega). \quad (13)$$

The equivalence of (8) and (13), together with the bound (12), shows that the best approximation in \mathcal{Q} to the posterior yields the tightest lower bound on the marginal likelihood.

We work with (13) rather than (8) – to evaluate the KL divergence, we would need to compute the marginal likelihood, bringing us back to our original problem. The variational parameters to be adjusted are $\lambda = \{\mu_{1:H}, \Sigma_{1:H}\}$. We conduct the variational inference (13) using block coordinate ascent on the coordinate blocks $\mu_1, \Sigma_1, \dots, \mu_H, \Sigma_H$. The coordinate updates do not have a closed form, but each update solves a smooth, unconstrained

convex optimization problem, as we show in Appendix A. There we also give a closed-form gradient and Hessian for the μ_h update, as well as closed-form gradients for the Σ_h update under two different parametrizations.

To finish with empirical Bayes, we explain how to obtain approximate MLEs $\hat{\zeta}$ and $\hat{\Omega}$. Notice the variational inference procedure (13) yields an optimal value λ^* for fixed ζ and Ω . We alternate this variational inference step with a complementary optimization over ζ and Ω . In fact, these optimizations constitute a version of the expectation-maximization (EM) algorithm for computing MLEs. The standard E-step, where we compute the posterior expected complete log likelihood, is replaced with a variational E-step, where the expected complete log likelihood is approximated using $q(\beta_{1:H} \mid \lambda^*)$. The variational EM algorithm alternates between the following steps until $\hat{\zeta}$, $\hat{\Omega}$, and the variational parameters stabilize:

E-step (variational). Using the current $\hat{\zeta}$ and $\hat{\Omega}$, run the block coordinate ascent algorithm as described in Appendix A, yielding new variational parameter values $\{\mu_{1:H}^*, \Sigma_{1:H}^*\}$.

M-step. Using the current variational parameter values, update the empirical Bayes parameter estimates: $(\hat{\zeta}, \hat{\Omega}) \leftarrow \operatorname{argmax}_{\zeta, \Omega} \mathcal{L}(\mu_{1:H}^*, \Sigma_{1:H}^*; \zeta, \Omega)$.

It can be shown that the variational E-step finds a new value of λ which moves the lower-bounding function $\mathcal{L}(\lambda; \zeta, \Omega)$ towards the true log-likelihood $\log p(\mathcal{D} \mid \zeta, \Omega)$ from below. The M-step maximizes this adjusted lower-bounding function over ζ and Ω , as a surrogate for the true log-likelihood. We then re-tighten and re-maximize until convergence. Appendix A gives details on initialization and the M-step update.

3.2 Variational hierarchical Bayes

Fully Bayesian inference in the mixed multinomial logit model requires calculations under the posterior $p(\beta_{1:H}, \zeta, \Omega \mid \mathcal{D})$ given in (7). In one sense, the setting is simpler than empirical Bayes: there are no unknown top-level parameters to estimate. All we need is to extend the previous section’s variational inference procedure to include ζ and Ω . Appendix A.4

reports the details behind the extension; here we summarize the main ideas.

Although the joint posterior (7) is not factorized, we continue to use a family \mathcal{Q} of factorized distributions for the variational approximation:

$$\mathcal{Q} \ni q(\beta_{1:H}, \zeta, \Omega \mid \lambda) := q(\zeta \mid \mu_\zeta, \Sigma_\zeta) q(\Omega \mid \Upsilon^{-1}, \omega) \prod_{h=1}^H q(\beta_h \mid \mu_h, \Sigma_h). \quad (14)$$

Using a factored family for a non-factored posterior is commonly called *mean-field* variational inference. In (14), $q(\zeta \mid \mu_\zeta, \Sigma_\zeta)$ is a K -variate normal density; $q(\Omega \mid \Upsilon^{-1}, \omega)$ is an inverse Wishart density; and the $q(\beta_h)$ factors are K -variate normal densities as before. In the analysis and the algorithm, it is convenient to use a well-known equivalence, treating $q(\Omega \mid \Upsilon^{-1}, \omega)$ as a Wishart distribution $\mathcal{W}(\Upsilon, \omega)$ on Ω^{-1} . We are therefore optimizing over variational parameters $\lambda := \{\mu_\zeta, \Sigma_\zeta, \Upsilon, \omega, \mu_{1:H}, \Sigma_{1:H}\}$. The variational problem for hierarchical Bayes is to find the best approximating distribution $q(\beta_{1:H}, \zeta, \Omega \mid \lambda^*)$ in \mathcal{Q} by solving the analog of (13):

$$\lambda^* = \operatorname{argmax}_{\lambda \in \Lambda} \mathcal{L}(\lambda). \quad (15)$$

As with empirical Bayes, we use a block coordinate ascent optimization algorithm to solve (15), iterating through the coordinate blocks that define λ . Here again, all coordinate updates are convex optimizations. The details appear in Appendix A: updates for μ_ζ , Σ_ζ , Υ all have simple closed forms; ω has a closed form which requires no updating; and the μ_h and Σ_h updates are similar to the empirical Bayes case.

4 Empirical results

We compared the accuracy and speed of the variational methods described in the previous section to a standard and widely used MCMC approach (Allenby and Rossi 2003), on a suite of simulated data sets. Each data set was generated using the discrete choice model given by (3) and (4). To simulate a data set with J choice items, K item attributes, and

H agents, we first fixed values of ζ and Ω , the parameters controlling the distribution of preferences in the agent population. We then independently drew a β_h vector for each agent, according to (4). We also drew for each agent 25 iid $J \times K$ item attribute matrices x_{ht} consisting of iid $N(0, 0.5^2)$ entries. Finally, for each agent, we used x_{ht} and β_h to simulate 25 choice events y_{ht} , according to (3). Thus, in our data sets, each agent has 25 observed choices.

We simulated a total of 32 different scenarios by varying J , K , H , and the selection of ζ and Ω . Specifically, each data set corresponds to a distinct configuration of the following candidate values: 3 or 12 choice items J ; 3 or 10 item attributes K ; 250, 1000, 5000, or 25000 agents H ; and “low” or “high” heterogeneity of the agent population. In the low-heterogeneity scenario, the $K \times 1$ vector ζ consists of evenly spaced values from -2 to 2, and the $K \times K$ matrix Ω is 0.25 times the identity matrix. In the high-heterogeneity scenario, ζ is the same, but Ω is the identity matrix. The data sets with high heterogeneity have much more diverse collections of preference vectors β_h .

We ran variational empirical Bayes (VEB), variational hierarchical Bayes (VB), and the standard MCMC algorithm on the observable data from each of the 32 simulation scenarios. For VEB, we declared convergence as soon as an E-step/M-step iteration caused the parameter estimates’ joint Euclidean norm to change by less than 10^{-4} , relative to their norm at the beginning of the iteration (here we mean the joint norm of all the variational parameters, together with the model parameter estimates $\hat{\zeta}$ and $\hat{\Omega}$). The convergence criterion for VB was the same, except ζ and Ω do not have point estimates – we look instead at the change in the variational parameters corresponding to their posterior approximation, namely μ_ζ , Σ_ζ , and Υ .

Choosing MCMC convergence criteria is more delicate. We tried to set the number of burn-in iterations and the thinning ratio algorithmically, using the technique of Raftery and Lewis (1992). But on several of our data sets, typical control parameter values, such as

the default settings for the `raftery.diag` function in the R package `coda` (Plummer et al. 2006), led to a very large number of burn-in iterations. Trace plots of the sampled parameters indicated these large burn-in values were unnecessary, so using them would have been unfair to MCMC in our timing comparisons. Instead, we manually investigated MCMC convergence and autocorrelation for several data sets, using trace plots, partial autocorrelation functions, and related diagnostics available in the `coda` package. Based on these studies, we chose to use 1,000 iterations of burn-in and a 10:1 thinning ratio. On each data set, we therefore ran 6,000 total iterations of MCMC to generate 500 draws for the approximate posterior. These numbers are as small as we could reasonably make them, in order to be fair to MCMC in the timing comparisons.

4.1 Accuracy

Our measure of accuracy for each inference procedure is based on the *predictive choice distribution*. Informally, this distribution gives the item choice probabilities exhibited by the “average” agent, when shown an item attribute matrix x_{new} . In our simulations, we know the true values of ζ and Ω , so we can compute the true predictive choice distribution:

$$p(y_{\text{new}} \mid x_{\text{new}}, \zeta, \Omega) = \int p(y_{\text{new}} \mid x_{\text{new}}, \beta) p(\beta \mid \zeta, \Omega) d\beta \quad (16)$$

$$= \mathbb{E}_{\beta} p(y_{\text{new}} \mid x_{\text{new}}, \beta) . \quad (17)$$

Equation (17) explains the “average agent” interpretation of the predictive choice distribution. A slightly different take on (16)-(17) is the following: if we want to forecast the item choice probabilities of a new, previously unobserved decision-maker, we can think of her as a random draw from the agent population. Under our model, the choice probabilities for such a randomly drawn agent are precisely (16)-(17).

For any particular x_{new} , we use the true values of ζ and Ω to compute a Monte Carlo estimate of (17). We base the estimate on enough draws of $\beta \sim \mathcal{N}(\zeta, \Omega)$ to insure that

its variability does not affect even the least significant digit of our reported results. We handle the integral over β in the same way for the estimated predictive choice distributions furnished by each of the three inference procedures. For VEB, the estimate is (16), with ζ and Ω replaced by $\hat{\zeta}_{\text{VEB}}$ and $\hat{\Omega}_{\text{VEB}}$. On the other hand, with VB and MCMC, we obtain a posterior distribution over ζ and Ω ; we take the mean of (16) under this posterior as a point estimate of the predictive choice distribution:

$$\hat{p}(y_{\text{new}} | x_{\text{new}}, \mathcal{D}) = \int \left[\int p(y_{\text{new}} | x_{\text{new}}, \beta) p(\beta | \zeta, \Omega) d\beta \right] p(\zeta, \Omega | \mathcal{D}) d\zeta d\Omega . \quad (18)$$

For VB, the posterior density $p(\zeta, \Omega | \mathcal{D})$ in (18) is approximated by the fitted variational distribution

$$q(\zeta | \mu_{\zeta}, \Sigma_{\zeta}) q(\Omega | \Upsilon^{-1}, \omega) . \quad (19)$$

For MCMC, the posterior is approximated as usual by the empirical distribution of draws from a Markov chain. In both cases, we handle the integral over ζ and Ω in (18) with another exhaustive Monte Carlo approximation.

We measure the error of each inference procedure as the distance from its estimate of the predictive choice distribution to the true distribution. As the metric on distributions, we use total variation (TV) distance, leading to what we call the “TV error” of each procedure. In this setting, TV error equals the maximum, over all choice-item subsets, of the difference in the probabilities assigned to the subset by the estimated versus the true choice distribution. We also need to choose the attribute matrix x_{new} at which the true and estimated predictive choice distributions are calculated. In each simulation scenario, we computed the TV error of VEB on 25 different random draws of x_{new} . We then compared the three procedures using the x_{new} which yielded the median of the 25 TV errors. In this sense our results are representative of accuracy for a “typical” item attribute matrix. However, results using any one of the 25 matrices were qualitatively the same in the examples we checked.

Agents	Het.	3 attributes			10 attributes		
		VEB	VB	MCMC	VEB	VB	MCMC
250	Low	0.23	0.31	0.35	0.66	0.50	0.59
	High	0.49	0.48	0.92	0.68	0.51	1.46
1,000	Low	0.74	0.63	0.46	0.49	0.17	0.14
	High	0.53	0.33	0.09	0.30	0.39	0.18
5,000	Low	0.25	0.37	0.35	0.56	0.59	0.07
	High	0.19	0.25	0.39	0.35	0.20	0.38
25,000	Low	0.63	0.74	NA	0.59	0.55	NA
	High	0.53	0.53	NA	1.60	1.10	NA

Table 1: Total variation error in percentage points, for simulated data sets with three choice items. MCMC results are unavailable in the 25,000 agent case because the sampler exhausted memory resources before converging. See the text for the definition of total variation error.

TV error values for the three inference procedures are presented in Table 1 for the 3-item simulation scenarios, and in Table 2 for the 12-item scenarios. The main conclusion we draw from Tables 1 and 2 is simple: on these data sets, there are no practical differences in accuracy among VEB, VB, and MCMC. The scale of the TV error for all the procedures is the same; that scale is larger in the 12-item case than the 3-item case, but all three procedures exhibit high accuracy on all data sets. The magnitude of our simulation study makes it difficult to carry out the replications required to put standard errors in these tables. But even if the differences among TV errors in every scenario were “significant” under a suitable definition, the patternless alternation in the identity of the most accurate method would make more specific conclusions dubious.

Figure 1 shows in a different way the comparable accuracy of the three procedures. When there are three choice items, any predictive choice distribution can be plotted as a point in the triangular simplex. The figure shows the close proximity of each procedure’s estimated choice distribution to the true distribution in one simulation scenario. Plots for all the other three-item scenarios are qualitatively the same. Figure 1 also shows contours of the VB and MCMC approximate posterior distributions. We see that VB is producing

Agents	Het.	3 attributes			10 attributes		
		VEB	VB	MCMC	VEB	VB	MCMC
250	Low	2.88	2.80	2.64	1.97	1.91	2.44
	High	1.44	1.94	1.64	2.43	2.37	2.62
1,000	Low	1.11	1.27	1.09	0.99	1.00	1.60
	High	0.98	1.18	1.18	1.99	2.05	1.96
5,000	Low	1.25	1.45	1.18	0.95	1.13	0.97
	High	1.14	0.98	1.10	0.71	0.91	0.92
25,000	Low	0.22	0.33	NA	0.51	0.53	NA
	High	0.99	0.57	NA	1.23	0.96	NA

Table 2: Total variation error in percentage points, for simulated data sets with 12 choice items. See the caption accompanying Table 1.




Figure 1: Triangle plot of the true predictive choice distribution and its estimates in the three-item case, for the simulation with 250 decision-makers, 10 attributes, and high heterogeneity. See accompanying text.

not only a posterior mean similar to MCMC, but also a similar posterior density in the neighborhood of the mean.

4.2 Speed

For each simulated data set, we ran the three procedures in turn on the same unloaded machine: a 64-bit dual-core 3.2 GHz Intel Xeon processor with 8 GB of main memory. For the MCMC inference, we used the `rhierMnlRwMixture` function in the R package `bayesm` (Rossi and McCulloch 2007), which has efficient vectorized implementations of the inner sampling routines. This package stores all MCMC draws in memory, however. For our largest data sets, with 25,000 decision-makers, the machine’s memory was exhausted before MCMC converged. We were able to run MCMC for 1,000 iterations in this case, which allowed us to extrapolate accurately the time that would have been required for 6,000 iterations. We implemented the variational algorithms in R, with compiled C components for the numerical optimization routines.

Figure 2 displays time to convergence on each data set for the three procedures, according to the convergence criteria previously described. Within each panel, convergence time is plotted as a function of the number of agents, for fixed values of the other simulation parameters. Note that the vertical axis shows convergence time on a logarithmic scale, to ease comparison of MCMC to the variational methods. All the procedures scale roughly linearly with the number of agents, which leads to the logarithmic curves seen in the figure. The conclusions are the same in all the scenarios we simulated: variational methods converge faster than MCMC, and the magnitude of the difference increases with the number of observed agents. MCMC uses two days of computation time for 25,000 agents with 12 choice items, 10 item attributes, and high heterogeneity, versus an hour for each of the variational techniques. In the same setting but with low heterogeneity, MCMC’s two-day computation compares with two hours for VEB and six hours for VB. In other scenarios, variational run times are measured in minutes, as opposed to hours or days for MCMC.




Figure 2: Timing results for variational empirical Bayes (VEB), variational hierarchical Bayes (VB), and MCMC. Within each panel, convergence time is plotted on the log scale as a function of the number of agents, for fixed values of the other simulation parameters (shown at the top of each panel).

5 Discussion

Variational methods allow estimation of hierarchical discrete choice models in a small fraction of the time required for MCMC. They open Bayesian inference to a wide class of problems for which resource constraints make MCMC intractable, such as the MML model with many heterogeneous units. For now, variational methods appear to be the only viable option in these cases.

Of course, one can use variational methods to estimate many more types of models than the MML model examined here. Within the MML family, it would be straightforward to add a utility scaling parameter, or allow the heterogeneous coefficients themselves to depend on observed covariates. The value of the variational approach is greatest when subsampling of data is ill-advised. For example, consider a linear model with heterogeneous coefficients on exogenous and endogenous covariates, where the available instrumental variables only weakly explain the endogenous part. To draw inferences about the covariances of the heterogeneous parameters, we may need a large amount of data to achieve reasonable power in hypothesis testing. MCMC is untenable here, but variational methods have promise. When factorized variational distributions are inadequate, alternatives such as mixtures of normals or Dirichlet processes (Blei and Jordan 2006) can be applied.

We emphasize that we do not advocate abandoning MCMC in favor of variational methods. On the contrary, we suggest using MCMC when possible. MCMC offers consistency guarantees with no analog in variational methods, and it has withstood decades of scrutiny. But we advise against subsampling the data (i.e., throwing out information), or discarding key modeling elements, simply to make the problem fit within the time and resource constraints of MCMC. The possibility of applying variational methods to previously intractable problems makes them an important addition to the statistician’s toolkit.

A Variational inference and parameter estimation

In this appendix we describe the variational inference and estimation procedures for the mixed multinomial logit model. A few words on notation: $A \succ 0$ means the matrix A is positive definite; $A \succeq 0$ means positive semidefinite; $|A|$ is the determinant of A . For a scalar function $s : \mathbb{R} \rightarrow \mathbb{R}$ and a vector $v \in \mathbb{R}^n$, $s(v)$ means the $n \times 1$ vector $(s(v_1), \dots, s(v_n))^\top$.

A.1 The empirical Bayes ELBO

For empirical Bayes in the MML model, the ELBO objective function (11) becomes

$$H(q) + \sum_{h=1}^H \mathbb{E}_q \log p(\beta_h \mid \zeta, \Omega) + \sum_{h=1}^H \sum_{t=1}^{T_h} \mathbb{E}_q \log p(y_{ht} \mid x_{ht}, \beta_h) . \quad (20)$$

The first term in (20) is the Shannon entropy of the variational distribution. The second and third terms are (minus) an unnormalized cross entropy – the missing normalization constant is the marginal likelihood. Recall that the variational distribution $q(\beta_{1:H} \mid \mu_{1:H}, \Sigma_{1:H})$ is a product of normal distributions $\mathcal{N}(\mu_h, \Sigma_h)$.

The first and second terms of (20) are straightforward to derive; the third term requires more attention. Using the multinomial logit mass function

$$p(y_{ht} \mid x_{ht}, \beta_h) = \prod_{j=1}^J \left[\frac{\exp(x_{htj}^\top \beta_h)}{\sum_{j'} \exp(x_{htj'}^\top \beta_h)} \right]^{y_{ht}^j} , \quad (21)$$

the third term becomes

$$\sum_{h=1}^H \sum_{t=1}^{T_h} \left[\sum_{j=1}^J y_{ht}^j (x_{htj}^\top \mu_h) - \mathbb{E}_{q_h} \log \left(\sum_{j=1}^J \exp(x_{htj}^\top \beta_h) \right) \right] . \quad (22)$$

The expected log-sum-exp in (22) has no closed form. For variational inference, we there-

fore approximate the ELBO objective function \mathcal{L} using a new objective function $\tilde{\mathcal{L}}$. To construct $\tilde{\mathcal{L}}$, we consider two alternatives: the zeroth-order and first-order delta method for moments (Bickel and Doksum 2007), which we call D0 and D1 respectively. D0 is equivalent to applying Jensen’s inequality to the expected log-sum-exp, resulting in the following lower bound to (22):

$$[\mathbf{D0}] \quad \sum_{h=1}^H \sum_{t=1}^{T_h} \left[\sum_{j=1}^J y_{ht}^j (x_{htj}^\top \mu_h) - \log \left(\sum_{j=1}^J \exp(x_{htj}^\top \mu_h + (1/2) x_{htj}^\top \Sigma_h x_{htj}) \right) \right]. \quad (23)$$

Here we used the usual formula for the mean of a lognormal random variable. In a different context, Blei and Lafferty (2007) consider an approximation equivalent to D0 but expressed using a redundant variational parameter.

For approximation D1, we restrict $\Sigma_h \succ 0$ to be diagonal, and define $\sigma_h := \log(\text{diag}\{\Sigma_h\}) \in \mathbb{R}^K$. Using results in Appendix B, we obtain the following approximation to (22):

$$[\mathbf{D1}] \quad \sum_{h=1}^H \sum_{t=1}^{T_h} \left[\sum_{j=1}^J y_{ht}^j (x_{htj}^\top \mu_h) - \log \left(\sum_{j=1}^J \exp(x_{htj}^\top \mu_h) \right) - \frac{1}{2} \exp(\sigma_h)^\top \Theta(\mu_h) \right], \quad (24)$$

with $\Theta(\mu_h) \in \mathbb{R}^K$ as defined in Appendix B. Notice that, unlike D0, approximation D1 does not preserve the guarantee that the optimal value of the variational optimization lower bounds the marginal likelihood. However, in our simulations, using D1 resulted in more accurate variational approximations to the posterior.

In this appendix we give a derivation based on approximation D0. The derivation for D1 is similar, but simpler, because Σ_h is treated as diagonal. Under D0, the final empirical

Bayes objective function is

$$\begin{aligned}
\tilde{\mathcal{L}}(\mu_{1:H}, \Sigma_{1:H}; \zeta, \Omega) = & \frac{1}{2} \sum_{h=1}^H \log [(2\pi e)^K |\Sigma_h|] \\
& - \frac{H}{2} \log ((2\pi)^K |\Omega|) - \frac{1}{2} \text{tr} \left[\Omega^{-1} \sum_{h=1}^H \{ \Sigma_h + (\mu_h - \zeta)(\mu_h - \zeta)^\top \} \right] \\
& + \sum_{h=1}^H \sum_{t=1}^{T_h} \left[\sum_{j=1}^J y_{ht}^j (x_{htj}^\top \mu_h) - \log \left(\sum_{j=1}^J \exp(x_{htj}^\top \mu_h + (1/2)x_{htj}^\top \Sigma_h x_{htj}) \right) \right].
\end{aligned} \tag{25}$$

The first line in (25) uses the well-known entropy of the normal distribution. The second line uses the cross-entropy of two normal distributions, also well known. The third line is approximation D0.

A.2 Empirical Bayes variational E-step

Here we describe a block coordinate ascent algorithm to maximize (25) over the variational parameters $\mu_{1:H}$ and $\Sigma_{1:H}$. Although the problem is not jointly convex in all these parameters, each μ_h and Σ_h coordinate update solves a smooth, unconstrained convex optimization problem. The requirement $\Sigma_h \succeq 0$ is satisfied after each update. We initialize the variational parameters at the maximum likelihood estimates from a homogeneous model (in which all agents share a common β value).

The concavity of (25) in μ_h follows from the fact that $\Omega \succ 0$ and from the convexity of the log-sum-exp function. We update μ_h using standard algorithms for unconstrained convex optimization (Boyd and Vandenberghe 2004), supplying an analytic gradient and Hessian as follows. Define the function $w(\mu, \Sigma, x)$ taking values in \mathbb{R}^J , with j th component $\exp(x_j^\top \mu + (1/2)x_j^\top \Sigma x_j)$, and normalized to sum to one across j . The gradient of $\tilde{\mathcal{L}}$ with respect to μ_h can then be written

$$\frac{\partial \tilde{\mathcal{L}}}{\partial \mu_h} = -\Omega^{-1}(\mu_h - \zeta) + \sum_{t=1}^{T_h} \sum_{j=1}^J \left[y_{ht}^j - w^j(\mu_h, \Sigma_h, x_{ht}) \right] x_{htj}. \tag{26}$$

Note the similarity of this gradient to the gradient from an L_2 -regularized multiple logistic regression: it consists of a contribution from the regularizer (the left-hand term), plus a residual-weighted sum of covariate vectors. Abbreviating $w(\mu_h, \Sigma_h, x_{ht})$ to w_{ht} , an argument using matrix differentials (Magnus and Neudecker 2007) gives the Hessian

$$\frac{\partial \tilde{\mathcal{L}}}{\partial \mu_h \partial \mu_h^\top} = -\Omega^{-1} - \sum_{t=1}^{T_h} [x_{ht}^\top \text{diag}\{w_{ht}\} x_{ht} - (x_{ht}^\top w_{ht})(x_{ht}^\top w_{ht})^\top] . \quad (27)$$

The Σ_h coordinate update is harder, because we need to insure that $\Sigma_h \succeq 0$. Using a reformulation, we can avoid making the constraint explicit, which would complicate the optimization. Let $\Sigma_h = L_h L_h^\top$ for a lower-triangular matrix L_h . Since $\Sigma_h \succeq 0$, one such L_h always exists—the Cholesky factor. We replace each Σ_h in $\tilde{\mathcal{L}}$ with $\Sigma_h(L_h) := L_h L_h^\top$, and optimize over the unconstrained set of lower-triangular matrices L_h .

The objective function (25) remains concave in L_h . To see this, compare the terms depending on $\Sigma_h = L_h L_h^\top$ to the function studied in Appendix C. We now give the gradient with respect to L_h . Standard matrix differentiation of (25) leads to the Σ_h gradient

$$\frac{\partial \tilde{\mathcal{L}}}{\partial \Sigma_h} = \frac{1}{2} \left[\Sigma_h^{-1} - \Omega^{-1} - \sum_{t=1}^{T_h} x_{ht}^\top \text{diag}\{w_{ht}\} x_{ht} \right] . \quad (28)$$

Again using matrix differentials and the Cauchy invariance rule, it is not hard to show that the gradient with respect to L_h is

$$\frac{\partial \tilde{\mathcal{L}}}{\partial L_h} = 2 \left(\frac{\partial \tilde{\mathcal{L}}}{\partial \Sigma_h} \right) L_h = L_h^{-\top} - \left(\Omega^{-1} + \sum_{t=1}^{T_h} x_{ht}^\top \text{diag}\{w_{ht}\} x_{ht} \right) L_h . \quad (29)$$

Note that this is the gradient with respect to a dense matrix L_h . Since we optimize over lower-triangular matrices, i.e. $\text{vech}(L_h)$, we need only use the lower triangular of the gradient. This is convenient for the term $L_h^{-\top}$: it is upper-triangular, so its lower triangle is a diagonal matrix. Furthermore, from a standard result of linear algebra, the diagonal entries

are simply $1/\ell_{ii}$, where the ℓ_{ii} 's form the diagonal of L_h .

In practice we do the μ_h and Σ_h updates in a single step by optimizing jointly over μ_h and L_h , which remains a convex problem.

A.3 Empirical Bayes M-step

In the M-step, we maximize (25) over ζ and Ω . Identifying the terms which depend on ζ , we recognize the usual Gaussian mean estimation problem. Further, (25) is easily seen to be concave in Ω^{-1} , with a closed-form solution of the corresponding first-order condition. We obtain the M-step updates

$$\hat{\zeta} \leftarrow \frac{1}{H} \sum_{h=1}^H \mu_h, \quad \hat{\Omega} \leftarrow \frac{1}{H} \sum_{h=1}^H \Sigma_h + \widehat{\text{Cov}}(\mu.). \quad (30)$$

Here $\widehat{\text{Cov}}(\mu.)$ is the empirical covariance of the μ_h vectors.

A.4 Variational hierarchical Bayes

In the fully Bayesian MML model, ζ and Ω have prior distributions, with corresponding variational factors given in (14). The ELBO in this case has the same form as (20), with two differences. First, $H(q)$ contains two new terms

$$H(q(\zeta \mid \mu_\zeta, \Sigma_\zeta)) + H(q(\Omega \mid \Upsilon^{-1}, \omega)). \quad (31)$$

Second, there are two new cross-entropy terms

$$\mathbb{E}_q \log p(\zeta \mid \beta_0, \Omega_0) + \mathbb{E}_q \log p(\Omega \mid S, \nu). \quad (32)$$

Also, the middle term of (20) changes in the fully Bayesian case, because ζ and Ω are now averaged over rather than treated as constants.

Using known formulas for normal and Wishart entropies, the two new entropy terms

are seen to equal

$$\frac{1}{2} \log [(2\pi e)^K |\Sigma_\zeta|] - \frac{\omega - K - 1}{2} D(\omega, \Upsilon) + \frac{\omega K}{2} + A_\omega(\Upsilon). \quad (33)$$

Here we used the expected log determinant of a Wishart random matrix

$$D(\omega, \Upsilon) := \log (2^K |\Upsilon|) + \sum_{i=1}^K \Psi \left(\frac{\omega + 1 - i}{2} \right) \quad (34)$$

and the log normalization constant of the Wishart distribution

$$A_\omega(\Upsilon) := \log \left[2^{\omega K/2} \pi^{K(K-1)/4} \prod_{i=1}^K \Gamma \left(\frac{\omega + 1 - i}{2} \right) \right] + \frac{\omega}{2} \log |\Upsilon| \quad (35)$$

(see, for example, Beal 2003). The new cross entropy terms for ζ and Ω work out to

$$-\frac{1}{2} \left\{ \log [(2\pi)^K |\Omega_0|] + \text{tr} \left(\Omega_0^{-1} [\Sigma_\zeta + (\mu_\zeta - \beta_0)(\mu_\zeta - \beta_0)^\top] \right) \right\} \quad (36)$$

and

$$-A_v(S^{-1}) + \frac{v - K - 1}{2} D(\omega, \Upsilon) - \frac{\omega}{2} \text{tr}(S^{-1} \Upsilon) \quad (37)$$

respectively. The middle term of (20) eventually becomes

$$-\frac{H}{2} \{K \log(2\pi) - D(\omega, \Upsilon)\} - \frac{\omega}{2} \text{tr} \left[\Upsilon \left(H \Sigma_\zeta + \sum_{h=1}^H (\Sigma_h + (\mu_\zeta - \mu_h)(\mu_\zeta - \mu_h)^\top) \right) \right]. \quad (38)$$

With these changes, it is not hard to see that $\tilde{\mathcal{L}}$ is concave separately in μ_ζ and Σ_ζ . The

first-order conditions for block coordinate ascent lead to the updates

$$\mu_\zeta \leftarrow (\Omega_0^{-1} + H\omega\Upsilon)^{-1} \left(\Omega_0^{-1}\beta_0 + \omega\Upsilon \sum_{h=1}^H \mu_h \right), \quad (39)$$

$$\Sigma_\zeta \leftarrow (\Omega_0^{-1} + H\omega\Upsilon)^{-1}. \quad (40)$$

By inspection, $\Sigma_\zeta \succeq 0$, so this constraint need not be explicitly enforced. Note the similarity to conjugate posterior updating: on the precision scale, Σ_ζ is the sum of the prior precision matrix Ω_0^{-1} and H copies of the variational posterior mean $\omega\Upsilon$ for Ω^{-1} . Similarly, μ_ζ is a precision-weighted convex combination of the prior vector β_0 and the empirical average of the variational posterior means $\mu_{1:H}$ for $\beta_{1:H}$.

The updates for Υ and ω are similarly straightforward to derive; we obtain

$$\omega \leftarrow \nu + H, \quad (41)$$

$$\Upsilon \leftarrow \left(S^{-1} + \sum_{h=1}^H (\Sigma_h + (\mu_\zeta - \mu_h)(\mu_\zeta - \mu_h)^\top) + H\Sigma_\zeta \right)^{-1}. \quad (42)$$

Notice that the solution (41) for ω involves only the constants ν and H . We compute ω once in advance, leaving it unchanged during the variational optimization.

B An application of the delta method

Let $f(v)$ be a function from \mathbb{R}^K to \mathbb{R} . According to the multivariate delta method for moments (Bickel and Doksum 2007),

$$\mathbb{E}f(V) \approx f(\mathbb{E}V) + \frac{1}{2} \text{tr} \left[\left(\frac{\partial f(\mathbb{E}V)}{\partial v \partial v^\top} \right) \text{Cov}(V) \right]. \quad (43)$$

Consider the case

$$f(v) = \log(1^\top \exp(xv)) , \quad (44)$$

where x is a $J \times K$ matrix whose rows are the vectors x_j^\top . Let $V \sim \mathcal{N}_K(\mu, \Sigma)$, and restrict Σ such that $\Sigma = \text{diag}\{\exp(\sigma)\}$ for $\sigma \in \mathbb{R}^K$. We can now rewrite (43):

$$\mathbb{E} \log (1^\top \exp(xV)) \approx \log (1^\top \exp(x\mu)) + \frac{1}{2} \Theta(\mu)^\top \exp(\sigma), \quad (45)$$

where $\Theta(\mu)$ is the diagonal of the Hessian of f , evaluated at the point μ . Define $s = 1^\top \exp(x\mu)$. Using matrix differentials, it can be shown that

$$\Theta(\mu) = s^{-1} (x \odot x)^\top \exp(x\mu) - s^{-2} (x^\top \exp(x\mu)) \odot (x^\top \exp(x\mu)), \quad (46)$$

where \odot denotes the Hadamard product.

To use the approximation (45) in an optimization over μ , we need to compute the gradient. The formula for $\Theta(\mu)$ makes this a more extensive but still mechanical exercise in differentials. One obtains

$$\begin{aligned} \frac{\partial}{\partial \mu} = s^{-1} x^\top e^{x\mu} + \frac{1}{2} x^\top \left[\left(s^{-1} \text{diag}\{e^{x\mu}\} - s^{-2} e^{x\mu} (e^{x\mu})^\top \right) (x \odot x) + \right. \\ \left. 2 \left(s^{-3} e^{x\mu} \{ (x^\top e^{x\mu}) \odot (x^\top e^{x\mu}) \}^\top - s^{-2} \text{diag}\{e^{x\mu}\} x \text{diag}\{x^\top e^{x\mu}\} \right) \right] \exp(\sigma). \end{aligned} \quad (47)$$

C A convexity result

Let a_1, \dots, a_d be scalars, c_1, \dots, c_d be n -vectors, $p, r > 0$, and $Q \succeq 0$. We show here that the function

$$f(B) = r \log |BB^\top| - p \text{tr}(QBB^\top) - \log \left(\sum_{j=1}^d \exp \{ a_j + c_j^\top B B^\top c_j \} \right) \quad (48)$$

is concave on the set of full-rank $n \times n$ matrices.

We argue that each of the three constituent terms, from left to right, is concave. The

second differential of $g(B) = r \log|BB^T|$ is

$$d^2g = d \operatorname{tr} \left\{ 2r B^{-1} dB \right\} = \operatorname{tr} \left\{ -2r \left[B^{-1} (dB) \right]^2 \right\}. \quad (49)$$

By Theorem 10.6.1 of (Magnus and Neudecker 2007), the Hessian of g is

$-2r K_n (B^{-\top} \otimes B^{-1})$, where K_n is the order- n commutation matrix and \otimes denotes the Kronecker product. We now show that $K_n (B^{-\top} \otimes B^{-1})$ is (matrix) positive-definite.

$$(d \operatorname{vec} X)^\top K_n (B^{-\top} \otimes B^{-1}) (d \operatorname{vec} X) = (d \operatorname{vec} X)^\top K_n \operatorname{vec} \left\{ B^{-1} (dX) B^{-1} \right\} \quad (50)$$

$$= (\operatorname{vec} dX)^\top \operatorname{vec} \left\{ B^{-\top} (dX)^\top B^{-\top} \right\} \quad (51)$$

$$= \operatorname{tr} \left\{ \left(B^{-1} dX \right)^2 \right\} \geq 0. \quad (52)$$

Equation (50) follows from the well-known fact that $\operatorname{vec} ABC = (C^\top \otimes A) \operatorname{vec} B$. Thus, the Hessian of g is negative definite, and $r \log|BB^T|$ is concave.

Concavity of the middle term in (48) follows in the usual way from the univariate convexity of the function

$$g(t) := \operatorname{tr} \left(Q(M + tP)(M + tP)^\top \right) = \sum_{i=1}^n (m_i + tp_i)^\top Q(m_i + tp_i) \quad (53)$$

for fixed matrices M and P , with columns m_i and p_i . To see that the rightmost term in (48) is concave, define

$$g_j(t) := a_j + c_j^\top (M + tQ)(M + tQ)^\top c_j$$

for $j = 1, \dots, d$ and fixed matrices M and Q . Each g_j is convex in t , and the rightmost term in (48) is (minus) the log-sum-exp function composed with the g_j 's. Concavity of this term in t , and hence in B , follows from (Boyd and Vandenberghe 2004; p. 86).

References

- Albert, James H., and Siddhartha Chib. 1993. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 88(422):669–679.
- Allenby, Greg M., and Peter J. Lenk. 1994. Modeling household purchase behavior with logistic normal regression. *Journal of the American Statistical Association* 89(428):1218–1231.
- Allenby, Greg M., and Peter E. Rossi. 1999. Marketing models of consumer heterogeneity. *Journal of Econometrics* 89(1):57–78.
- . 2003. Bayesian statistics and marketing. *Marketing Science* 22(3):304–328.
- Beal, Matthew J. 2003. Variational algorithms for approximate Bayesian inference. Ph.D. thesis, University College London.
- Ben-Akiva, Moshe, and Steven Lerman. 1985. *Discrete choice analysis: Theory and application to travel demand*. Cambridge, Mass.: The MIT Press.
- Bickel, Peter J., and Kjell A. Doksum. 2007. *Mathematical statistics: Basic ideas and selected topics*, vol. 1. 2nd ed. Upper Saddle River, N.J.: Pearson Prentice Hall.
- Blei, David, and Michael I. Jordan. 2006. Variational inference for Dirichlet process mixtures. *Bayesian Analysis* 1(1):121–144.
- Blei, David, and John D. Lafferty. 2007. A correlated topic model of Science. *Annals of Applied Statistics* 1(1):17–35.
- Boyd, Stephen, and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge University Press.
- Bradlow, Eric T., and Elaine L. Zanutto. 2006. Data pruning in consumer choice models. *Quantitative Marketing and Economics* 4(3):267–287.
- Fader, Peter S., and Bruce G. S. Hardie. 1996. Modeling consumer choice among SKUs. *Journal of Marketing Research* 33(4):442–452.
- Guadagni, Peter M., and John D. C. Little. 1983. A logit model of brand choice calibrated on scanner data. *Marketing Science* 2(3):203–238.
- Hall, Jane, Patricia Kenny, Madeleine King, Jordan Louviere, Rosalie Viney, and Angela Yeoh. 2002. Using stated preference discrete choice modelling to evaluate the introduction of varicella vaccination. *Health Economics* 11:457–465.
- Jordan, Michael I., Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence Saul. 1999. An introduction to variational methods for graphical models. *Machine Learning* 37:183–233.

- Magnus, Jan R., and Heinz Neudecker. 2007. *Matrix differential calculus with applications in statistics and econometrics*. 3rd ed. Wiley Series in Probability and Statistics, Wiley.
- McFadden, Daniel L. 1974. The measurement of urban travel demand. *Journal of Public Economics* 3:303–328.
- Moore, William L., Jordan Louviere, and Rohit Verma. 1999. Using conjoint analysis to help design product platforms. *Journal of Product Innovation Management* 16(1):27–39.
- Plummer, Martyn, Nicky Best, Kate Cowles, and Karen Vines. 2006. CODA: Convergence diagnosis and output analysis for MCMC. *R News* 6(1):7–11.
- Raftery, Adrian E., and Steven Lewis. 1992. How many iterations in the Gibbs sampler? In *Proceedings of the fifth Valencia international conference on Bayesian statistics*, ed. Jose M. Bernardo, James O. Berger, A. Philip Dawid, and Adrian F. M. Smith, 763–774. Oxford University Press.
- Revelt, David, and Kenneth E. Train. 1998. Mixed logit with repeated choices: Households' choices of appliance efficiency level. *Review of Economics and Statistics* 80(4): 647–657.
- Robbins, H. 1955. An empirical Bayes approach to statistics. In *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability*, 131–148.
- Rossi, Peter, and Robert McCulloch. 2007. bayesm: Bayesian inference for marketing/micro-econometrics. *R package version 2.1-3*.
- Rossi, Peter E., Greg M. Allenby, and Robert McCulloch. 2005. *Bayesian statistics and marketing*. Wiley Series in Probability and Statistics, Chichester, UK: John Wiley and Sons.
- Theil, Henri. 1969. A multinomial extension of the linear logit model. *International Economic Review* 10(3):251–259.
- Train, Kenneth E. 2003. *Discrete choice methods with simulation*. Cambridge: Cambridge University Press.
- Train, Kenneth E., Daniel L. McFadden, and Moshe Ben-Akiva. 1987. The demand for local telephone service: A fully discrete model of residential calling patterns and service choices. *The RAND Journal of Economics* 18(1):109–123.
- Uhler, R. S., and J. G. Cragg. 1971. The structure of asset portfolios of households. *Review of Economic Studies* 38:341–357.
- Wainwright, Martin J., and Michael I. Jordan. 2003. Graphical models, exponential families, and variational inference. Tech. Rep. 649, UC Berkeley Statistics.